SPALDING

File Number 7010-00
Form A22-6724

IBM Systems Reference Library

IBM7010 Systems Summary

Brief descriptions are provided of the system concepts and special features, the system units and I/O devices, and the programming systems and available programs. Summarizing the fundamental facts about the 7010 system, this publication has the purpose of helping the user achieve a basic understanding of the system and the interrelationships of its many parts.

Publications providing detailed information on subjects discussed here are listed in the IBM 7010 Bibliography, Form A22-6720.

Address comments regarding this publication to:
IBM Corporation, Customer Manuals, Dept. 298, PO Box 390, Poughkeepsie, N.Y.

CONTENTS

IBM 7010 Data Processing System

The IBM 7010 Data Processing System handles problems and data volumes in magnitudes that characterize the intermediate to large scale data processing area. Outstanding capabilities and features combined with a wide systems configuration make the 7010 system a powerful general purpose data processing system. In addition, the 7010 is very effective in real time telecommunication applications and system-to-system communication. A wide variety of input and output devices allow the system to be readily expanded for growth in both applications and volume of work.

Important features and capabilities of the 7010 system are:

1. Two-character access (2.4 microsecond) core storage in sizes up to 100,000 character storage positions.
2. Powerful instruction repertoire, including high-speed arithmetic operations, 64 different move operations, editing, table lookup and other special operations augmented by indexing and chaining functions.
3. Up to two data channels with buffered input-output operations, including high-speed printing, card reading, and punching.
4. Magnetic tape processing at as many as eight data rates. Tape data exchange with the system at rates up to 180,000 characters per second performed concurrently with processing.
5. Large capacity disk storage (28 to 280 million characters) providing large volume random access storage for in-line data processing to frequently used tables, to programmed subroutines, or to other data.
6. Standard features of process overlap and priority processing.
7. Special checking features to assure reliability and accuracy of results.
8. A complete library of programs and programming systems including Autocoder, input-output control systems, sort-merge and a variety of utility programs to simplify program testing and systems operation.
9. Complete instruction compatibility with the IBM 1410 Data Processing System. The 7010 is designed such that programs produced by IBM programming systems for the 1410, operate unaltered when input-output configurations are identical, with no distinction either in input or results.

The 7010 system can include a variety of versatile and fast I/O devices. Cards, for example, are read at speeds as high as 800 cards per minute, and punched at 250 cards per minute. The 1403 Printer prints at the rate of up to 600 lines per minute. For large volumes of numeric printing, the alphameric chain can be replaced with a numeric chain to more than double the printing speed. The 1403 accomplishes fast document throughput owing to the high speed of line skipping (75 inches per second).

Ten 729 or 7330 Magnetic Tape Units can be attached to each data channel of the 7010. Four models of the 729 (Models II, IV, V, VI) are available for use. With the tape intermix feature, 729's and 7330's can be intermixed on a channel. Direct data exchange with the system can occur at rates up to 180,000 characters per second, performed concurrently with processing.

In addition to being the primary on-line input or output medium of the 7010 system, tape can also be the means of exchanging data between a 7010 and a peripheral processor, such as a 1410 or 1401 system. Efficiency and effectiveness of the 7010 increase substantially when another system acts as a peripheral data converter (card-to-tape, tape-to-card, or tape-to-printer) and editor.

The large capacity, high-speed random access storage afforded by 1301 Disk Storage units can be combined with the powerful internal capabilities of the 7010, to provide an efficient and economical in-line processing system. A 1301 Disk Storage Unit can contain as many as 56 million characters of random access storage. Up to five disk storage units can be used with the system. Data exchange with the 7010 can be performed at a rate of 90,000 characters per second. Records varying in length up to the total length of a track can be stored in any sequence in a track. Because any record on any track is addressable, the 7010 has direct access to any record in the 1301. This random accessability is the key to the in-line approach to data processing; it eliminates the necessity for accumulating transactions of like kind (batching) before entering them. Transactions can be entered as they occur, regardless of other defined logical sequence.

A 1301 can be simultaneously on-line to two systems and therefore can be directly shared between a 7010 and another 7010 or between a 7010 and a 1410 or 7000 series computer (except the 7072 system). There are three general methods of sharing a disk storage between two systems:

1. Two systems share a common file of information (programs or data).
2. Two systems exchange data; one system places the information on disk, the other system uses the information.

3. Each system uses the file at different times or uses different areas of the file independently, as temporary or as permanent storage.

Many present-day data processing applications require that data operated upon by a central computer be exchanged within short time intervals with data processing equipment remote from the central system. A configuration of data processing equipment that provides for centralized data processing in a physically decentralized environment and uses telecommunications facilities for data transmission, is known as a Telecommunication system.

The data processing equipment at the remote locations may consist of a number of terminal devices (e.g., IBM 1013, IBM 7702, IBM 1014 or Teletype* terminals) or other IBM computer systems. In addition to the telecommunications attachments, the central processor can still service the normal input-output units such as card reader, magnetic tape units, disk files, printers and punches, as required. 7010 systems with telecommunications attachments generally perform in-line/on-line operations, processing data as it originates and updating immediately all records affected. Operations using single batch or schedule-type data may also be performed.

The 7010 system has two general systems configurations that include telecommunications attachments. The specific configuration is determined by the requirements of the application. One such system configuration utilizes the IBM 7750 Programmed Transmission Control for medium to large telecommunication capability.

The IBM 7750 is a Tele-Processing® component that links the 7010 system to vast telecommunication facilities and remote data processing equipment. The 7750 operates under control of its own stored program. It performs automatic and programmed functions such as assembly and distribution of messages, transfer of data to and from the IBM 7114 Processing Unit, deletion and indication of certain communication and process control errors, code conversion, editing, validity checks, count keeping, message numbering, and in some cases even format changes.

The 7750 performs the varied functions of telecommunication in such a manner that the 7010 system, instead of dealing with the many individual elements of the telecommunication facilities and remote data processing equipment, now deals with only one unit (the 7750), which appears as simply another input-output device. In effect, the 7750 greatly simplifies the job of real time processing.

Services provided by the 7750 include:

1. IBM 7701 and 7702 Magnetic Tape Transmission Terminals
2. IBM 1009 Data Transmission Units
3. IBM 1013 Card Transmission Terminals
4. IBM 65/66 Data Transceiver
5. Standard Telegraph Terminals

The 7010 system requiring small to moderate telecommunication capability can effectively utilize the facilities provided by the IBM 1414 Models 4 and 5 Input-Output Synchronizers. These I/O Synchronizers provide on-line attachment of communication oriented and paper tape devices such as:

1. IBM 1009 Data Transmission Unit
2. IBM 1014 Remote Inquiry Unit
3. Telegraph Input and Output Units
4. IBM 1011 Paper Tape Reader

The 1009 Data Transmission Unit provides the 7010 system with efficient and economical two-way data transmission between two remote 7010 systems, between a 7010 and an IBM 7701 or 7702 Magnetic Tape Converter (or IBM 1013 Card Transmission Terminal) or between a 7010 and another IBM system (1400-7000 series) equipped with a 1009. Data can be transmitted in either direction over leased or toll lines.

The IBM 1014 Remote Inquiry Unit, using typewriter input and output, is another means of system interrogation, providing a visual record of information entered into and transmitted from the system. Remote inquiry provides direct access to any record stored within the 7010 system and furnishes printed output under program control. Remote inquiry units can be located up to eight wire miles away from the system.

Common carrier equipment may also be attached to serve as remote input-output devices. As many as four telegraph adapters may be attached to a 1414 Model 4 or 5 to communicate with remote input-output telegraph units. The data transmission rate of these attachments can be up to approximately 10 characters per second, depending on the transmission rate of the common carrier equipment used. Direct input from these devices represents a further expansion of in-line processing.

The flexibility of the 7010 system is further enhanced by the availability of the IBM 1011 Paper Tape Reader for direct paper tape input. Paper tape is an ideal medium for recording information because it is available from a variety of source recording devices. In addition, information on paper tape can be transmitted quickly and inexpensively by telecommunications equipment at the central location.

*Trademark of Teletype Corporation

The basic 7010 system consists of the 7114 Model 1 Processing Unit with 40,000 positions of core storage, the 7010 instruction repertoire: priority processing and overlap, and data channel 1 with its associated 1415 Console, Model 2. Optional input-output synchronizers and control units for channel 1 are:

1. IBM 1414 Input-Output Synchronizer, Model 1, 2, or 7.
2. IBM 1414 Input-Output Synchronizer, Model 3, 4, 5, or 8
3. IBM 7631 File Control, Models 1, 3, or 5
4. IBM 7750 Programmed Transmission Control, Models 1, 2, or 3

Optional features allow the expansion of core storage to 60,000, 80,000 or 100,000 character positions, addition of a second data channel (identical to channel 1 except for the 1415 Console, Model 2) and the 1401 compatibility feature.

The 7010 is available in such a variety of configurations that a system can be tailored with the units and features that exactly suit individual requirements --with the added advantage of permitting the system to be modified or expanded as needs change or increase. All possible combinations of units and features are shown in the IBM 7010 System Configuration. In addition, guides are available for planning an efficient and suitable physical installation. For the configuration and the separate manuals and bulletins on physical planning, see IBM Systems Reference Library, IBM 7010 Bibliography, Form A22-6720.

SYSTEM CONCEPTS

CORE STORAGE

The basic 7010 system is provided with a 40,000-character-position core storage unit. A 60,000, 80,000, or 100,000 character position storage unit is available as an optional feature. Each storage position is individually addressable with a unique five-character address. Addresses used to access storage are checked to ensure that they are not greater than the available storage size.

The basic read-write cycle time of core storage is 2.4 microseconds. Each storage access allows two adjacent characters (in parallel) to be placed in storage or removed from storage for processing. This double character access reduces total storage access time and thus increases the effective processing speed of the system. All control and alignment functions peculiar to the two character parallel storage are automatically performed. No particular areas of storage are reserved exclusively for the program. The location of instructions, constants, or data to be processed is entirely at the discretion of the programmer. Core storage is housed in the IBM 7114 Processing Unit.

CODING AND CHARACTER SET

Information stored and moved within the system is coded in an eight-bit binary coded decimal form (Figure 1). Six of the bits (B, A, 8, 4, 2, 1) are used as information bits to generate the standard 64 character set. The seventh bit (C) is used to maintain odd parity on all internal characters. The eighth bit is called a word mark bit (WM) and is primarily used to define the starting and/or ending positions of data fields and instruction fields. Word marks are set and cleared by instruction. The word mark bit is symbolized, in printed form, by an inverted circumflex (Â) over the character containing a word mark bit.

Each character is checked at various locations in the system to be sure that the number of bits, including the check bit and the word mark bit, is odd. If a parity error occurs, a console light turns on to indicate the error. The character set and collating sequence are identical to that employed by the 1410 system.

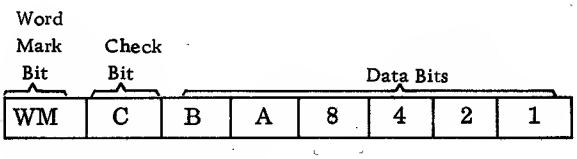

Figure 1. Character Code Bits

WORDS AND WORD MARKS

The 7010 is a variable word length system. Data and instructions are manipulated in groups of characters called words. This variable word length principle allows words to occupy only the number of spaces in storage that are actually needed for a specific operation.

Word length is controlled by the use of word marks. Word marks indicate the first character of an instruction and signal the end of an instruction. The length of a data field usually defined with word marks, although it may be defined with other symbols such as record marks. Essentially word marks make possible the variable length words of the 7010. They can be set and cleared as necessary by specific instructions of the stored program.

Word marks can be translated to word separator characters when written out of storage onto cards, tape or disk, and can be retranslated to word marks when read back into core storage. In this manner word marks can be stored on I/O devices.

INSTRUCTION WORD

The 7010 is a stored program system using variable-length double-address type instructions. Each instruction consists of a variable number of alphameric characters that combine to make up an instruction word. The actual operation performed is indicated by the format and contents of the instruction itself.

The basic instruction format is divided into four parts: operation code, the A-address or an x-control field, the B-address, and a d-character modifier to the operation code (Figure 2).

Part 1	Part 2	Part 3	Part 4
Op Code	A-Address or x-control field	B-Address	d-character
O	aaaaa or xxx	bbbbb	d

Figure 2. IBM 7010 Instruction Format

Operation Code is a single character that specifies the basic machine operation to be performed.

A-Address is a five-character core storage address. In some operations the A-address is identified as an I- or C-address.

x-Control Field, three characters in length, specifies the channel and I/O unit involved in an input-output operation.

B-Address is a five-character core storage address. d-Character is used to specify a particular operation within the control of the operation code of the instruction.

Valid instruction words vary in length from one to twelve characters depending on the amount of information required for the operation. Instruction formats consist of a single-character operation code, or the operation code with one or more of the other three parts. Valid instruction lengths and formats are:

O	Oxxxbbbbbd
Od	Oaaaaabbbbb
Oxxxd	Oaaaaabbbbbd
Oaaaaa	
Oaaaaad	

Instruction length checking is incorporated in the system to insure that each instruction read contains a valid number of characters for the operation code specified.

Each instruction must have a word mark set over the operation code, and must not contain word marks in any other position. Also a word mark must be set in the core storage location immediately to the right of the last character of an instruction. This is normally the word mark associated with the operation code of the next sequential instruction.

IBM 7010 instructions are the two-address type. That is, a single instruction can address two data fields or two different characters, and so on. In arithmetic operations, for example, one instruction addresses both factors, such as the addend and augend or the multiplier and multiplicand. Double addressing is also characteristic of data moving and most other operations, but instructions with single addresses are also used--in some cases optionally, and in others necessarily.

PROGRAM BRANCHING

Instructions are assigned sequential locations in storage. The method of reading and executing these instructions is also sequential, unless the sequence is altered by an instruction that causes a program branch to another storage area for the next instruction.

The branch may be unconditional. On the other hand, the program may make a conditional branch by examining conditions that arise during processing. If a comparison is equal, or if a specific character exists in a storage, or if any of a large number of other tests is satisfied, the program may transfer to a predetermined instruction or subroutine, or repeat a group of instructions.

CHAINING

Whenever, at the end of an operation, certain machine registers contain the addresses of the next data fields to be processed, these addresses need not be given in the next instruction. Because the registers already hold the addresses of the fields, the operation code alone can be given. The contents of the machine registers are used as the address portion of the instruction to be chained. Sequential instructions can frequently be chained in this manner; the gain is to save most of the storage space otherwise required for the next instruction, also to save most of the time needed by the machine to read that instruction.

INDEXING

Any address of any instruction (except the store address register instruction) may be indexed (modified) by the contents of one of fifteen five-digit index registers. Each index register uses five core storage positions, and its contents are called the index factor. This factor is added or subtracted, according to its sign, from any address that is tagged (with zone bits) for such modification. Thus, a single instruction or a series of instructions can have a continuously changing effect.

Seventy-five locations of storage (00025-00099) are assigned for use as index registers. These locations are addressable in the same manner as any other locations and are available as general storage if not required for indexing. Zone bits over the tens and hundreds position of the address in the instruction are used as tags to specify if the address is to be indexed and which index register is to be used. Addresses greater than available storage size may be used in instructions, provided that the address is indexed to a permissible size.

7010 OPERATIONS

A general description of certain basic operations of the 7010 system follows.

Arithmetic

The physical make-up of the system makes it possible for the IBM 7010 to do arithmetic directly in the storage area. (This is called add-to-storage logic.) The use of add-to-storage logic in the 7010 system eliminates the need for special purpose accumulators or counters in the system. Because any group of storage positions can be used as an accumulating field, the capacity of arithmetic quantities is not limited by a finite number of counter or accumulator positions.

Example: Add first field to second field.

Instruction: A aaaaa bbbbb

Each address locates the units position of its field. Numerical data in the field located by the A-address are added to the numerical data in the field located by the B-address. The sum is stored in the B-field. The A-field is in no way changed.

The same form of instruction is used for subtract, multiply, and divide. The main difference is in the operation code.

Data Moving

Whole data characters consist of a numeric portion, a zone portion, and a word-mark portion. Data moving causes a copy of any selected portion(s) of all the characters in one storage area (A-field) to be moved to another area (B-field), where the copy combines with or replaces the data, if any, in the B-field. Data can be moved in 64 different ways. The criteria are the direction of the move, the kind of data moved, and the terminating condition.

The direction may be left to right or right to left; the kind may be numeric portion only, word marks only, zone portion only, or any combination of these, including a move of whole characters with word marks at one extreme and no portion at all (a scan) at the other. The amount of data moved can be a single position, or the move can be continuous until stopped by a word mark in the A or B field or the first record mark or group mark-word mark in the A field.

Editing

Editing punctuates and adds special characters in data lines before they are printed. Editing also suppresses unwanted leading zeros, and can close up the remaining significant quantity with dollar sign or asterisk protection as required. Credit symbols (CR) and minus signs can be indicated. One instruction loads into the print area a control mask containing commas, points, dollar signs, a 0 signifying zero suppression, etc., and a second instruction edits the raw data by distributing the data within the mask.

Comparing

The 7010 system compares data fields by testing the bit structure of each character in the A-field with that of each character in the B-field. The compare operation in no way changes either field, but produces one of three results: the B-field is higher than, lower than, or equal to the A-field. The result can be interrogated with a subsequent branch instruction. The compare is not algebraic but follows the collating sequence.

Table Lookup

A table lookup instruction causes the machine to search through a table for rates, mathematical factors, or other types of information. The information is called a function. The machine searches for this function by making a comparison between a data field (called a search argument) external to the table, and another data field (called a table argument) stored in the table. Table search stops when the desired kind of comparison is found. At this point the address of the function is held in an address register; the function thus located can now be used for processing.

The table has no size limit other than the capacity of core storage.

PROCESSING OVERLAP

Processing overlap allows computing to continue most of the time taken by the transfer of data to or from input-output devices and core storage. Input-output data move to and from core storage via a two-character channel buffer. From the channel buffer, data are exchanged with the buffers of the I/O devices attached to the system or, in the case of I/O devices not buffered, with the devices themselves. The time required to transfer two characters of data between core storage and the channel buffer is substantially less than the time required to transfer two characters of data between the channel buffers and an I/O buffer. During overlapped I/O operations, processing is interrupted only for the time required to transfer two characters to or from core storage via the channel buffer. In other words, processing continues while the two characters of data are exchanged between the channel buffer and the I/O buffer or I/O device. Also, processing may be performed during the time an I/O device is starting, stopping or otherwise preparing to send or receive data.

If a 7010 system has two data channels, reading and writing can be overlapped with processing, thus utilizing the full power and capabilities of the 7010 system.

PRIORITY PROCESSING

Priority processing is an interrupt system that provides an automatic interrupt branch to core storage location (00101) when certain conditions of the

I/O channels or devices occur. These conditions are stored in indicators which may be tested to determine the specific cause of the interrupt.

Priority processing provides increased efficiency in the use of system units. Virtually all the waiting time can be eliminated from any input-bound or output-bound points in a program. Processing is not held up until a device finishes a certain operation. A selected device signals the stored program as its operation is completed, and, on the basis of these signals, the program determines the sequence of steps needed to utilize best all units of the system. The devices can automatically interrupt the main routine by a request for service or by indicating that a certain condition exists. An interrupt transfers system control from the main routine to a priority routine. With the priority feature two independent programs can be processed, so that one is functioning while the other is waiting for one of its operations to end.

IBM 1401 COMPATIBILITY FEATURE

The design of the 7010 system permits running many programs as originally written for 1401 system. The 1401 compatibility feature (optional) allows the 7010 system to function like a 1401 system with 16K core storage and the following Advanced Programming features: multiply and divide, compare, expanded print edit, sense switches, and process overlap. Details about this feature can be found in the IBM Systems Reference Library, IBM 1401 Compatibility Feature for IBM 7010 System, Form A22-6721.

SYSTEM UNITS

This section expands the section "System Configurations and Physical Planning" by briefly describing the units that can be used to tailor an IBM 7010 Data Processing System for specific applications.

7114 PROCESSING UNIT

The 7114 Processing Unit contains the arithmetic and control circuitry of the system, as well as the various address registers, logic and checking circuits, etc., associated with the flow of data in the 7010. The 7114 is available in four models, depending on the quantity of core storage housed within the unit:

Model 1	-- 40,000 positions
Model 2	-- 60,000 positions
Model 3	-- 80,000 positions
Model 4	--100,000 positions

1415 MODEL 2 CONSOLE

The 1415-2 Console (Figure 3) is made up of a control section containing the keys that control the 7010 system, an indicator panel, and the system's I/O printer. The I/O printer types 64 characters, a word mark, and an underscore symbol, providing:

1. A log of all console operations done manually, including a required log of any manual change made in internal data and addresses.
2. A display of the contents of any storage location or addressable register.
3. Printed-out messages under program control (print-out can be overlapped with computing).
4. Print-out of the various address registers on any manual stop, programmed stop, or error stop.
5. An inquiry mode of operation, under which the operation can stop the program to change it or request information from it.

1414 INPUT/OUTPUT SYNCHRONIZER

The 1414 Input/Output Synchronizer (Figure 4) contains the circuitry necessary for transmitting data to and from the processing unit and I/O equipment. Figure 5 shows the 1414 models necessary to control the various kinds of I/O devices.

Model 3 of the 1414 houses the 80-character card read and punch buffers and a 100- or 132- character print buffer, with all associated controlling circuitry.

Model 4 is the same as Model 3, but has additional control facilities for telecommunication, including controls and buffers for the 1011 Paper Tape Reader, the 1009 Data Transmission Unit, as many as 20 Remote Inquiry Units, and the Telegraph Input/Output Feature. Model 5 is the same as Model 4 but lacks card read, card punch, and print buffers. The 1414 model 8 has a 100- or 132-character print buffer with associated controlling circuitry. It differs from Models 3 and 4 in that the attachment of the 1402 Card Read Punch is not required.

The tape intermix feature, and any others that are optional or required to complete an installation, are

Figure 3. IBM 1415-2 Console

Figure 4. IBM 1414 Input/Output Synchronizer

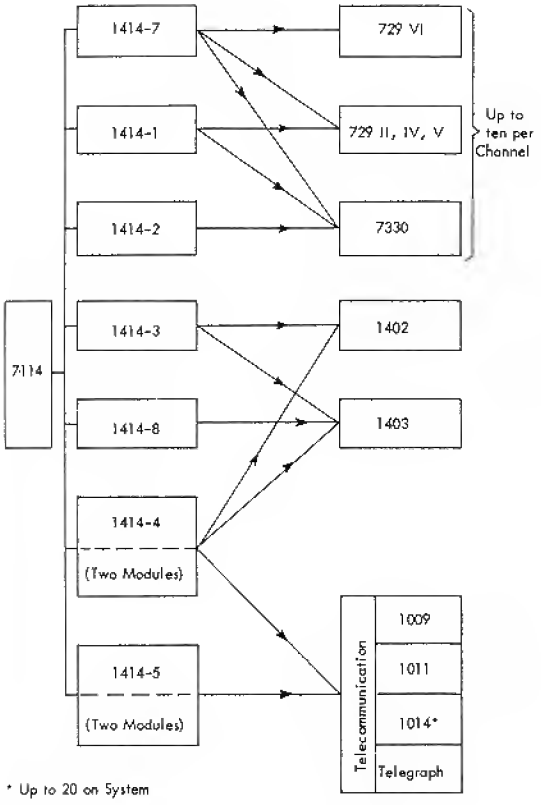

Figure 5. Synchronizer Models Required for Each I/O Device

shown in the IBM 7010 System Configuration. (See IBM Systems Reference Library, IBM 7010 Bibliography, Form A22-6720.)

1402 CARD READ PUNCH

The 7010 system uses Model 2 of the 1402 Card Read Punch (Figure 6), which can read and punch cards simultaneously at a read speed of 800 cards per minute and a punch speed of 250 cards per minute. The punch feed hopper has a capacity of 1,200 cards; the read file feed can be loaded with as many as 3,000 cards.

The reader enters data into the system via an 80-position read buffer in the 1414, and the punch receives data from the system via an 80-position punch buffer. The 1414 Input/Output Synchronizer allows the data channel of the 7010 to perform other I/O operations during card reading and punching operations.

Stacker

The card stacker has five pockets, with a capacity of 1,000 cards each. Cards read are stacked in pockets 2, 1, or 0 (normal read); cards punched are stacked in pockets 8, 4, or 0 (normal punch). The program controls reading, punching, and stacking, and determines the pocket selected.

51-Column Cards

Stub 51-column cards, commonly used for installment payments, postal money-order forms, inventory cards, etc., need not be reproduced on standard 80-column cards before entry into the system. An interchangeable read feed and file feed permits feeding either standard 80-column cards or stub 51-column cards directly into the data processing system. The operation of the 1402 is not affected by the interchangeability.

1403 PRINTER

The 1403 Printer, Model 1, produces a standard printed line 100 positions long. With the use of Model 2, it is possible to increase the length of the printed line to 132 positions (Figure 7). The typical alphameric chain can print 48 different characters in each position; these are the 26 letters of the alphabet, the 10 numbers, and 12 special characters such as * and $. Ten characters are printed per inch.

A dual-speed carriage, controlled by a closed loop of paper tape, permits skipping at 33 inches per second for skips of eight lines or less, or at 75 inches per second for skips of more than eight lines. Because document printing involves a large amount of

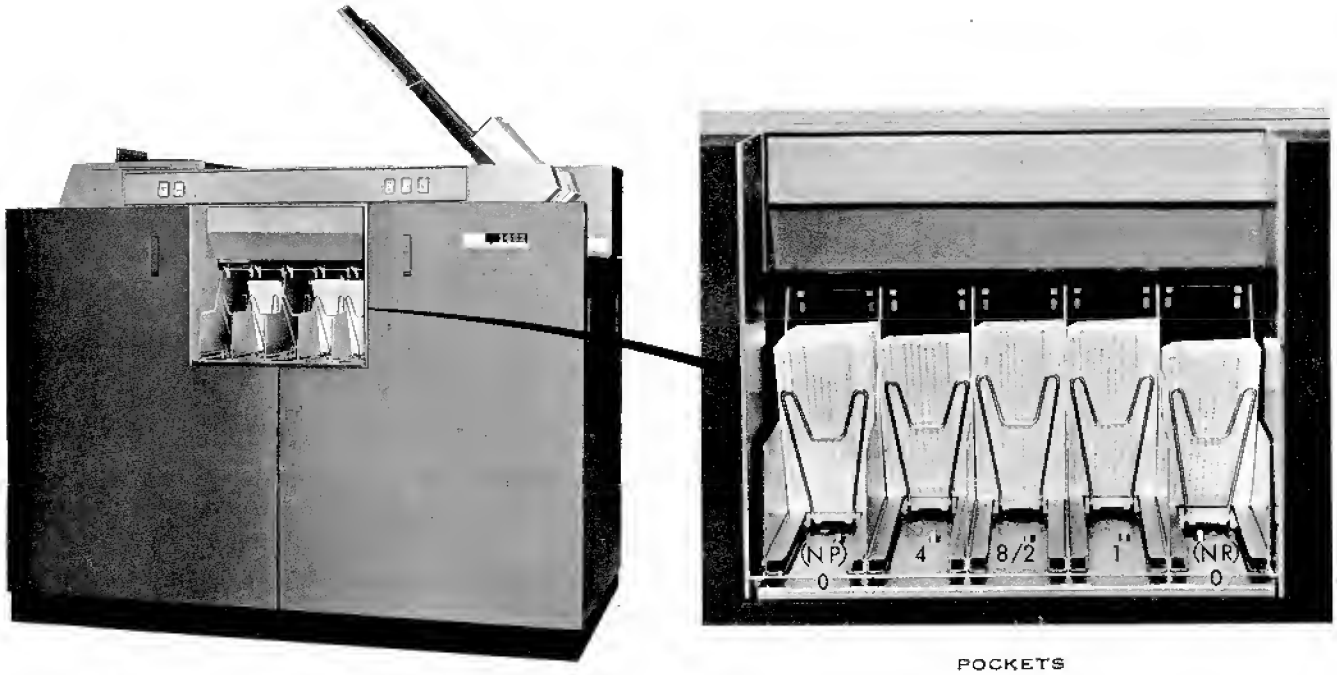

Figure 6. IBM 1402 Card Read Punch, Model 2

form skipping, this high-speed skip provides the fast document throughput that marks printer efficiency.

The standard alphameric chain permits printing at 600 lines per minute. The lines are spaced either six or eight to the inch, as selected by the user. Vertical form-movements, such as a space, a double space, or a skip, are under control of the stored program.

The printer receives data from the system via a 100- or 132-position print buffer in the 1414 Input/Output Synchronizer. The 1414 allows the data channel of the 7010 to perform other I/O operations during a printing operation.

Special Chains

The alphameric chain cartridge can be replaced with a numeric chain cartridge, if warranted by a large amount of numeric printing, with the gain of doubling the printing speed. The numeric chain contains only 16 characters (0 through 9 and $. , - * ⊟), but prints at a speed of 1285 lines per minute.

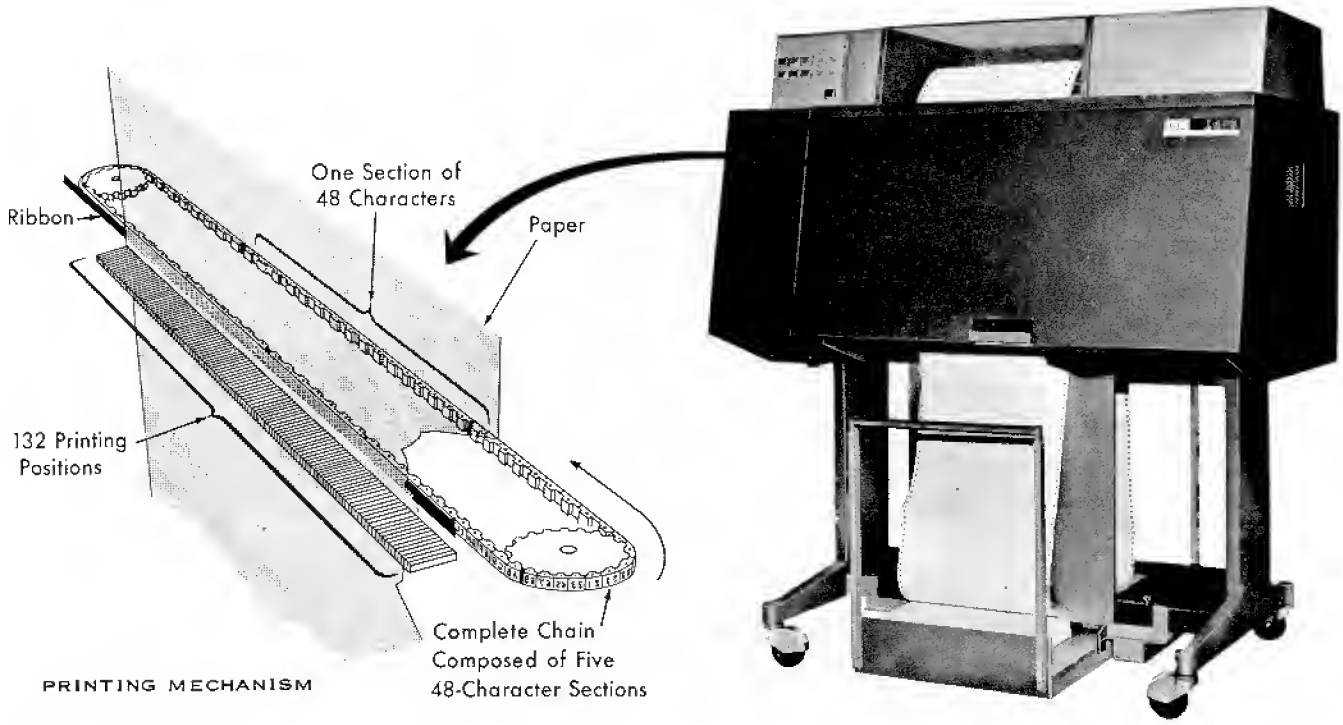

Figure 7. IBM 1403 Printer

Many commercial and scientific applications require distinctive type styles for particular printing jobs. Without using special tools, the operator can replace the chain cartridge with another containing type of a different font or style, or having a particular character arrangement.

Printing Method

All characters are serially assembled in a closed chain that revolves horizontally. A magnet-driven hammer taps the paper form against the chain as each character to be printed travels past an appropriate print position.

729 AND 7330 MAGNETIC TAPE UNITS

Ten 729 or 7330 Magnetic Tape Units (Figure 8) can be attached to each data channel of the 7010 system. Four models of the 729 (Model II, IV, V, or VI) are available for 7010 use.

The essential difference between the 729 and 7330 tape units is the character rate. Many work loads demand the use of the fastest 729's, but a customer having smaller volumes of data to process can have the advantages of a tape system at a reduced cost by using 7330's. With the tape intermix feature, moreover, 729's and 7330's can be intermixed on the same data channel.

Differences in character rate, character density, and tape speed are shown in Figure 9. The higher recording densities (characters per inch) provide the significant advantage of storage on fewer tape reels per given volume of data. All units have dual recording density, and thus are able to transfer tape from one IBM system to another. The 729 Models V and VI can be conditioned by a switch to operate as determined by the program at either the two highest or the two lowest character rates depicted by the bar graphs.

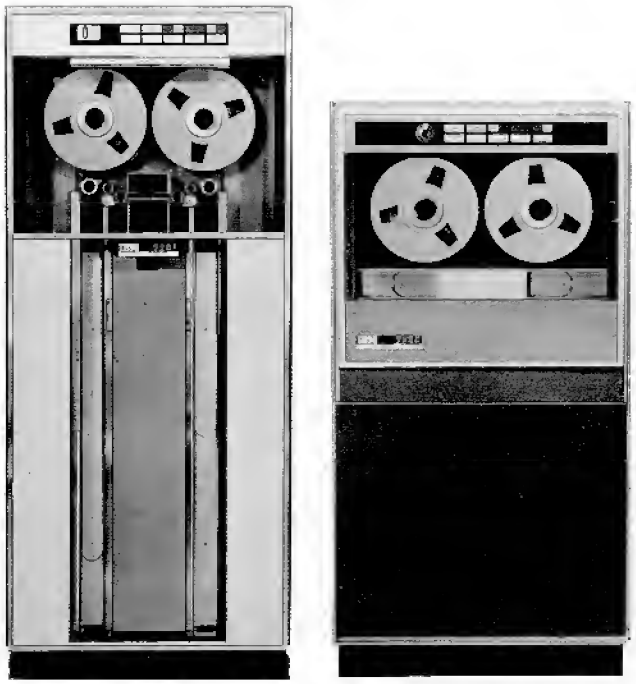

Figure 8. IBM 729 and 7330 Magnetic Tape Units

1301 DISK STORAGE UNITS

Five 1301 Disk Storage Units (Figure 10) can be attached to the 7010 system; the maximum number is five whether the particular system has one or two data channels. Total disk storage can vary from 28 to 280 million character positions.

The large-capacity, high-speed, random access storage of the 1301 combines with the data processing capabilities of the 7010 to provide an efficient and economical in-line system. In-line data processing keeps business records constantly up to date. Any transaction can be processed when it occurs, all records and accounts affected are immediately updated, and the status of any account at that moment is readily at hand. Two 1301 models are available:

Model 2 has two disk storage arrays and stores up to 56,000,000 alphameric characters in variable size records up to 2800 characters long. Each array has its own access mechanism.

Model 1 has one disk storage array, and thus half the storage capacity of Model 2.

The 1301 has a maximum character rate of 90,000 characters per second.

The 7631 File Control used for the 1301's has three models that are available for use with the 7010:

Model 1 is for 7010 use exclusively.

Model 3 is for shared use with the 7010 and any IBM 7000 series system (except another 7010 and the 7072).

Model 5 is for shared use with two 7010 systems or 7010 and 1410 systems.

The disk file controls, optional features, and adapters completing any disk storage installation are shown in the IBM 7010 System Configuration (see IBM Systems Reference Library, IBM 7010 Bibliography, Form A22-6720.)

1011 PAPER TAPE READER

One 1011 Paper Tape Reader (Figure 11) can be attached to a 7010 system, providing for data sources stored on paper tape. Data can be read at up to 500 characters per second. The paper tape can be in widths commonly used for 5-, 6-, 7-, or 8-track

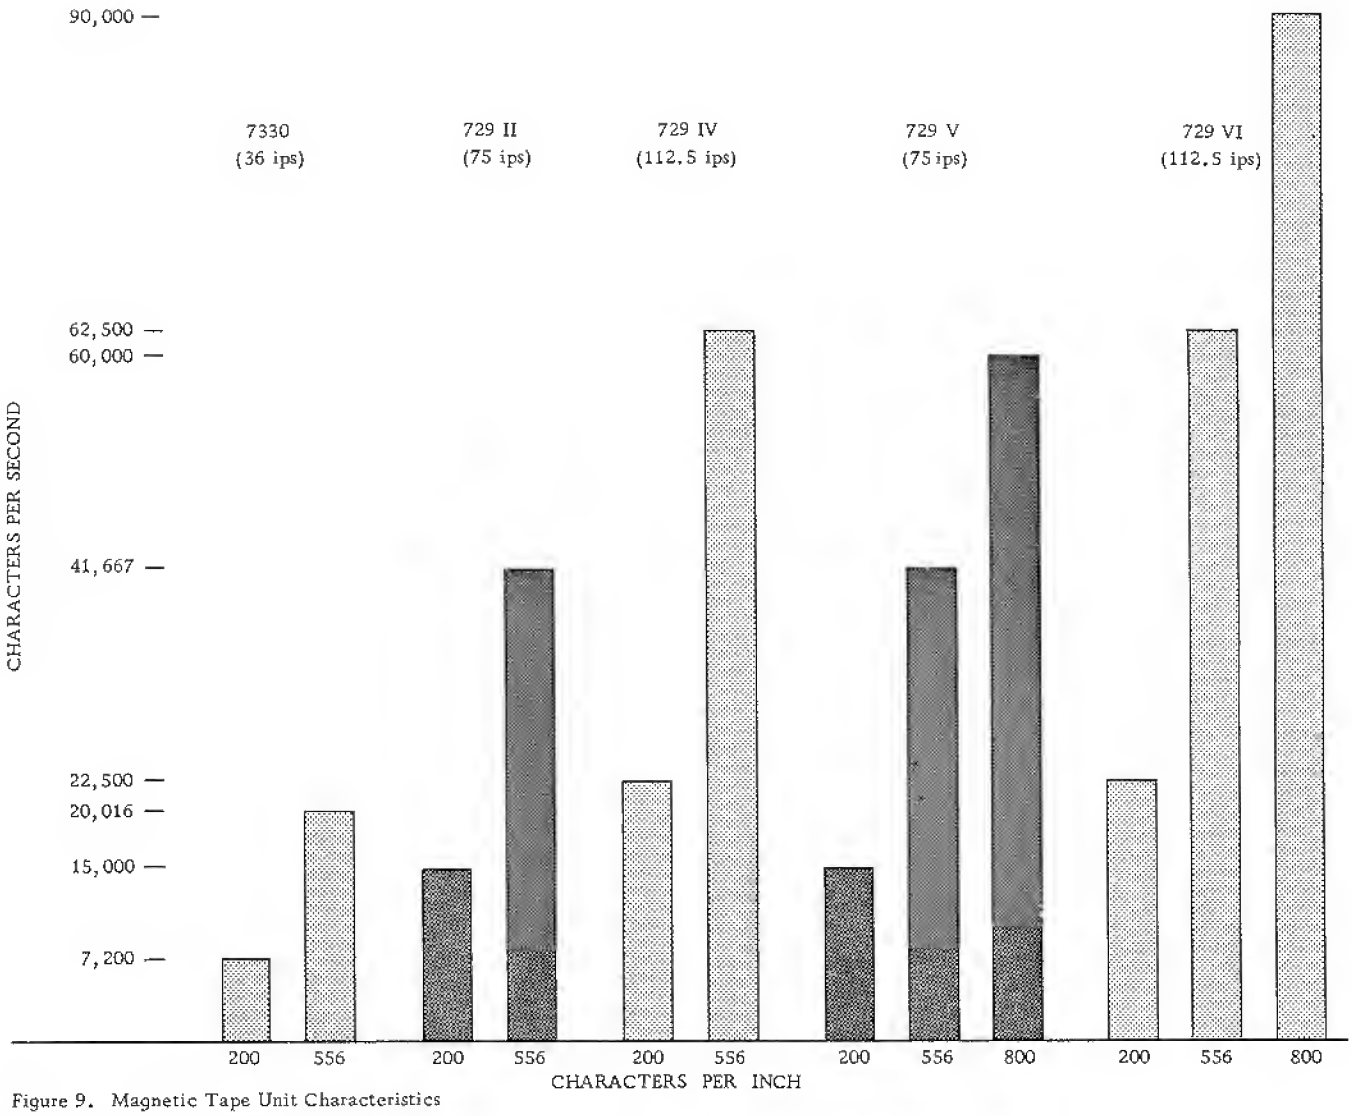

Figure 9. Magnetic Tape Unit Characteristics

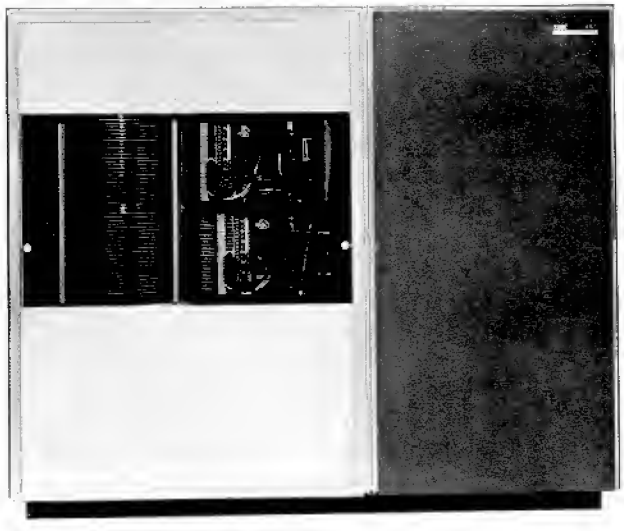

Figure 10. IBM 1301 Disk Storage Unit

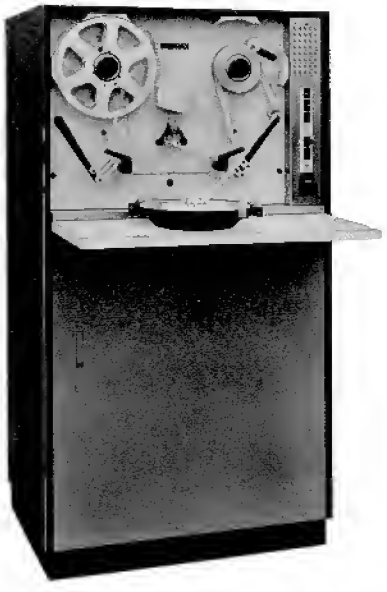

Figure 11. IBM 1011 Paper Tape Reader

tape, in strips, reels, or rolls that feed from the center, and can be chad or chadless.

Data are read from the 1011 via an 80-position buffer in Model 4 or 5 of the 1414 I/O Synchronizer. The 1414 allows the 7010 data channel to perform other I/O operations during a paper tape reading operation.

7750 PROGRAMMED TRANSMISSION CONTROL

The 7750 Programmed Transmission Control (Figure 12) serves as a buffer and telecommunications message control unit that links the 7010 system with a variety of remote transmitting and receiving terminal devices. The 7750's stored program assembles messages, distributes messages under priority to and from the network, converts codes, edits, checks messages for validity, etc., while directing data at high character rates to and from the host 7010.

One 7750 can be assigned to a system. Model numbers (1, 2, or 3) of the 7750 denote differences in the size of process storage. The 7750, Model 1 has 4096 word positions of process storage, Model 2 has 8,192 words and Model 3 has 16,384 words.

TELEGRAPH INPUT/OUTPUT

Installation of the Telegraph Input/Output Feature in Model 4 or 5 of a 1414 permits connecting a telegraph network to the 7010 system. Data can thus be accepted directly by the 7010 from any station in the network, or transmitted directly from the 7010, through local telegraph terminals, to any or all of the stations.

The feature consists of one input adapter and one output adapter in the 1414 I/O Synchronizer. One or two more adapters can be added; these additional adapters can consist of one input or one output adapter (or both), or two input or two output adapters. Data are transferred to and from the telegraph via 80-position buffers, one for each input or output adapter. The 1414 allows the 7010 data channel to perform other I/O operations during a transfer to or from the telegraph network.

Figure 12. IBM 7750 Programmed Transmission Control

1009 DATA TRANSMISSION UNIT

The 1009 Data Transmission Unit (Figure 13) is the intermedium for high-speed two-way communication between the 7010 system and another 1009 (or a 1009 simulated by a 7750 Programmed Transmission Control) attached to any 1400- or 7000-series data processing system. Data are transmitted or received over message-service or leased wire circuits at fixed character rates ranging from 75 to 300 per second.

Data are transferred to or from the 1009 via two 80-position buffers in Model 4 or 5 of the 1414 Input/Output Synchronizer. The 1414 allows the 7010 data channel to perform other I/O operations during a 1009 operation.

1014 REMOTE INQUIRY UNIT

A 7010 system containing one or more 1014 Remote Inquiry Units (Figure 14) can reply to remote or local requests for information. Each 1014, whether installed locally or remotely, receives and prints almost instantaneous replies from the 7010, and prints both the inquiry and the reply, at character rates up to 12-1/2 per second for the inquiry and 15-1/2 per second for the reply.

An adapter, in Model 4 or 5 of a 1414 I/O Synchronizer, controls and sequences the acceptance by the 7010 of request messages from as many as ten 1014's. Either one or two adapters can be installed in the 1414, to handle up to ten remote inquiry units each. With two adapters, two 1014's controlled by separate adapters can transmit or receive at the same time.

Inquiries and replies are transferred to and from the 1014 via two 80-position buffers, one for input and one for output, in each adapter. The 1414 also allows the 7010 data channel to perform other I/O operations during a 1014 operation.

Figure 13. IBM 1009 Data Transmission Unit

Figure 14. IBM 1014 Remote Inquiry Unit

These formulas give the execution times (in microseconds) for the various internal instructions. Since timings are dependent on the storage location of the instruction and the data fields, average timing formulas are given. The following symbols are used in the following formulas:

A	is the A-field length.
B	is the B-field length.
C	is 1 if the branch is taken. It is zero otherwise.
D	is the number of characters in the B-field from the start of zero suppression to the place where the dollar sign is inserted. If no dollar sign is inserted, D is zero.
E	is 2 on a single-character multiply or divide operation. It is 1 on a single-character add, subtract, reset add, reset subtract, table search, or a 6-character multiply or divide operation. It is zero otherwise.
I/O	is the time used by input/output device to accept or send data and the synchronizer access time, when applicable.
L	is the instruction length.
M	is the multiplier length.
N	is the number of fields actually compared on a table search. The B-field length on a table search includes only those argument fields actually compared and the intervening function values.
Q	is the quotient length.
R	is 1 if a re-complement is taken on an add or subtract operation. It is zero otherwise.
TM	tape movement
Z	is the number of characters in the B-field from the start of zero suppression (as indicated in the control field) to the left end of the B-field.

Add, Subtract, Reset Add, Reset Subtract (two fields)

$$T = 2.4\left(\frac{L+2}{2}\right) + 2.4E + 2.4\left(\frac{A+1}{2}\right) + 3.2\left(\frac{B+1}{2}\right) + 3.2\left(\frac{RB+1}{2}\right)$$

Add, Subtract, Reset Add, Reset Subtract (one field)

$$T = 2.4\left(\frac{L+2}{2}\right) + 5.6\left(\frac{A+1}{2}\right)$$

Multiply

$$T = 2.4\left(\frac{L+2}{2}\right) + (2.5\,M+1)\left[5.6\left(\frac{A+2}{2}\right) + 6.4\right]$$

Divide

$$T = 2.4\left(\frac{L+2}{2}\right) + 6.5\,Q \quad \left[5.6\left(\frac{A+2}{2}\right) + 6.4\right]$$

Indexing

$$T = 9.6 \text{ per address}$$

Branch Unconditional

$$T = 2.4\left(\frac{L+2}{2}\right)$$

Test and Branch (conditional)

$$T = 2.4\left(\frac{L+2}{2}\right)$$

Test and Branch if Channel Status Indicator On

$$T = 2.4\left(\frac{L+2}{2}\right)$$

Test Character and Branch

$$T = 2.4\left(\frac{L+4}{2}\right)$$

Test Bit and Branch

$$T = 2.4\left(\frac{L+4}{2}\right)$$

Test Zone or WM and Branch

$$T = 2.4\left(\frac{L+4}{2}\right)$$

Data Move

$$T = 2.4\left(\frac{L+2}{2}\right) + 2.4\left(\frac{A+1}{2}\right) + 3.2\left(\frac{B+1}{2}\right)$$

Move and Zero Suppress

$$T = 2.4\left(\frac{L+2}{2}\right) + 2.4\left(\frac{A+1}{2}\right) + 3.2\left(\frac{B+1}{2}\right)$$

Compare

$$T + 2.4\left(\frac{L+2}{2}\right) + 2.4\left(\frac{A+1}{2}\right) + 3.2\left(\frac{B+1}{2}\right)$$

Edit

$$T = 2.4\left(\frac{L+2}{2}\right) + 2.4\,A + 3.2\,B + 3.2\left(\frac{Z+1}{2}\right) + 3.2\left(\frac{D+1}{2}\right)$$

Table Lookup

$$T = 2.4\left(\frac{L+2}{2}\right) + 3.2\left(\frac{B+1}{2}\right) + 2.4\,N\left(\frac{A+1}{2}\right)$$

Store Address Register

$$T = 2.4\left(\frac{L+2}{2}\right) + 9.6$$

Store and Restore Status

$$T = 2.4\left(\frac{L+2}{2}\right) + 2.4$$

Set WM, Clear WM

$$T = 2.4\left(\frac{L+2}{2}\right) + 6.4$$

Clear

$$T = 2.4\left(\frac{L+2}{2}\right) + 3.2\left(\frac{B+1}{2}\right)$$

Clear and Branch

$$T = 2.4\left(\frac{L+2}{2}\right) + 3.2\left(\frac{B+1}{2}\right)$$

Halt

$$T = 4.8$$

Halt and Branch

$$T = 2.4\left(\frac{L+2}{2}\right)$$

No Op

$$T = 2.4\left(\frac{L+2}{2}\right)$$

I/O M, L, F, K, 2, 4, (except file)

$$T = 2.4\,(L+1) + I/O$$

I/O Unit Control

$$T = 2.4\,(L+1) + TM$$

I/O M, L (File)

$$T = 2.4\,(L+1) + \text{Address Transfer} + I/O$$

This section describes the programs and programming systems that will be initially available for the IBM 7010 Data Processing System. Because the IBM 7010 is compatible with the IBM 1410 Data Processing System, these programs and programming systems, which were originally written for the 1410, will operate on the 7010.

The programs described in this section are divided into "programming systems" and "service programs. Both concepts are defined below.

Purpose of the Programming Systems

Because writing programs in actual machine language is a laborious task, greatly subject to human error, "symbolic" languages have been developed to speed and simplify programming. These languages permit the programmer to represent storage locations and instructions by easily recognizable names and symbols. A particular advantage is the elimination of the need to pre-assign storage locations.

Such a language, together with a "processor program" that converts it into machine language, is called a programming system. Four programming systems will be available for the IBM 7010 Data Processing System: Autocoder, RPG (Report Program Generator), COBOL, and FORTRAN. These programming systems are incorporated in a "Processor Operating System."

In addition to the processors noted above, the Processor Operating System contains routines designed to control the flow of information between the computer and various input/output devices. Collectively, this set of routines is called the input/output control system. All or selected portions of these routines can be used to create a specific input/output control system for each machine-language program produced by the Processor Operating System.

Purpose of the Service Programs

In addition to the programming systems, IBM provides completely written "service programs." In contrasting with the programming systems, which simplify program writing, the service programs are finished products that perform specific functions under the guidance of control information supplied by the user of the programs.

Service programs that will be available for the IBM 7010 include Utility Programs, Sorting/Merging Programs, and a Simulation Program.

PROGRAMMING SYSTEMS

The Processor Operating System

The Processor Operating System (POS) is a set of programs designed to produce object programs from source programs written in the Autocoder, RPG (Report Program Generator), COBOL, and FORTRAN symbolic languages. This set of programs consists of the four programming-system processors, a group of "system supervisory programs," and a "system library."

The system supervisory programs, in addition to performing control functions, provide the facilities for updating and duplicating the POS, printing specified sections of the system library, and adding (to the library) macros and subroutines written by the customer's programmers.

The POS library consists of macros and subroutines, and includes the input/output control system.

There will be two versions of POS available for the 7010. One is designed to use magnetic tape for storage of the system and for intermediate processing; the other is designed to use 1301 Disk Storage.

The principal features of the individual systems included in POS are described as follows:

Autocoder

Autocoder is a symbolic language that can produce object programs for all 7010 machine configurations. The symbolic statements are punched into cards (one statement per card). These cards are called the "source program." The processor reads the source program, translates the symbolic statements into machine language, assembles operands, and assigns core-storage locations for each instruction, work area, and constant. After this translation and assembly is complete, the processor produces a machine-language program deck (the "object" deck). The object deck is preceded by a short program that can clear storage, set high-order word-marks in the index registers, and load the object program into storage. In addition, the processor produces a printed listing of the source statements and their machine-language equivalents. Coding errors are also indicated on this listing.

Report Program Generator

The Report Program Generator (RPG) is a programming system designed to creat object programs that

can write reports. The programmer used the RPG language to give certain control specifications to the processor program, and the processor produces a program that will read the input records and write the desired report.

Object programs produced by this programming system will read input records contained on cards, magnetic tape, and 1301 Disk Storage. Reports written by the object programs can be printed, punched in cards, and put on magnetic tape.

COBOL

COBOL (COmmon Business Oriented Language) is a programming system designed primarily for commercial data processing. It is the result of work by the Conference on Data Systems Languages (CODASYL), which is a voluntary cooperative effort by a number of users and manufacturers of data processing systems.

Writing in COBOL is similar to writing in English, and programmers with COBOL experience for one machine system can quickly learn to write in COBOL for other systems.

The COBOL processor, using macro-instructions and the library of the Processor Operating System, produces a symbolic program. The symbolic program is translated and assembled into a machine-language program by the Autocoder processor.

FORTRAN

FORTRAN (FORmula TRANSlating System) is a programming system designed primarily for scientific and technical applications. The FORTRAN language is similar to the language of mathematics. Because of this, personnel with mathematical backgrounds, but without programming training, can quickly learn to prepare problems for computers.

FORTRAN, like COBOL, is a language that can be adapted to different machine systems. This minimizes program rewriting and programmer retraining when an installation changes to a new computer.

The FORTRAN processor produces a symbolic program that is translated and assembled into a machine-language object program by the Autocoder processor.

Input/Output Control System

A significant portion of the instructions in an average program are related to input-output operations, error detection and correction, and end-of-file routines. The Input/Output Control System (IOCS), which is contained in the library of the Processor Operating System, provides pre-written routines to handle these functions.

The IOCS includes macro-instructions and routines for unit-record equipment, magnetic tape units, 1301 Disk Storage, and telecommunication devices connected through the 1414 Synchronizer. After the programmer defines his particular machine configuration, he can use the IOCS macro-instructions to generate (through the Autocoder processor) appropriate blocking, deblocking, and scheduling routines. He can also use labeling routines and checkpoint-and-restart routines provided by the IOCS.

SERVICE PROGRAMS

Utility Programs

Utility programs perform basic service functions that contribute a significant portion of daily machine operations. Besides operations such as storage prints and tape duplications, this includes utilities that function as program-testing aids. The brief descriptions below summarize the purpose of each utility program.

Card and Tape Utilities

Load. This program loads information punched in cards into core storage.

Clear Storage. This program clears storage to blanks.

Tape File Generators (A and B). These programs create tape files from punched cards.

Tape Print. This program produces a listing of information contained on magnetic tape.

Tape Compare. This program compares the information on two tapes and lists the records and record numbers that are not identical.

Tape Duplicate. This program transfers the information from one magnetic tape onto another.

Write-Tape-Mark-and-Rewind. This program writes a tape mark at the current position of a specified tape, and then rewinds the tape.

Storage Print. This program lists the contents of core storage.

Storage Punch. This program punches the contents of core storage into cards.

Snapshot. This is a program-testing aid that lists the contents of a selected area of core storage each time a specified instruction is executed in the object program.

Trace. This is a program-testing aid that lists all executed instructions (with relevant information) within a specified section of the object program.

Branch Trace. This is a program-testing aid that lists only executed branch instructions of the object program.

1301 Disk Utility Programs

Format Track Generation Program. This program writes one or more format tracks according to the specifications of the user.

Home Address and Record Address Generation Program. This program writes home addess identifiers and record addresses on one or more tracks in accordance with the user's specifications. It is always loaded into storage with the format track generation program. It can be executed only after the format tracks have been written.

Load Disk Program. This program loads information from magnetic tape into specified areas of disk storage.

Dump Disk Program. This program writes information from one or more tracks of disk storage onto magnetic tape.

Restore Disk Program. This program reloads into disk storage information that was written onto magnetic tape by the dump disk program.

Clear Disk Program. This program clears selected areas of disk storage and fills those areas with blanks or any other 7010 character specified by the user.

Sorting/Merging Programs

IBM provides generalized programs that can perform a wide variety of sorting and merging functions. These programs can process files of blocked and unblocked records that are of fixed length or variable length. Labeling and checkpoint-and-reset facilities are included, as are exit points for incorporation of specialized routines written by the user.

Sort/Merge 11

This program enables users to sort and merge files written on magnetic tape. Although it was designed for 1410 systems that do not have the Processing Overlap and Priority special features, it will operate on the 7010.

Sort/Merge 12

This program performs the same functions as Sort/Merge 11, but because Sort/Merge 12 was designed for 1410 systems that do not have the Processing Overlap and Priority special features, it will function more efficiently on the 7010 system than Sort/Merge 11.

Sorting Program Using IBM 1301 Disk Storage

This program is designed to use 1301 Disk Storage as working storage for sorting files. It will sort files contained on either 1301 Disk Storage or magnetic tape.

Simulation Program

To facilitate transition from an IBM 650 to an IBM 7010, a "simulation program," will be available for the 7010 system. Use of this program permits the 7010 to execute programs that were originally written for the 650. Thus, operations can be maintained during the time the programs are being rewritten for the 7010. Also, programs that are used infrequently, such as those for monthly and yearly reports, need not be rewritten at all.

Printed in U.S.A. A22-6724

IBM®
International Business Machines Corporation
Data Processing Division, 112 East Post Road, White Plains, New York